On Reward Function for Survival

Naoto Yoshida
Tohoku University, Sendai, Japan
E-mail: naotoyoshida@pfsl.mech.thoku.ac.jp

arXiv:1606.05767v2 [cs.AI] 24 Jul 2016

***Abstract*—Obtaining a survival strategy (policy) is one of the fundamental problems of biological agents. In this paper, we generalize the formulation of previous research related to the survival of an agent and we formulate the survival problem as a maximization of the multi-step survival probability in future time steps. We introduce a method for converting the maximization of multi-step survival probability into a classical reinforcement learning problem. Using this conversion, the reward function (negative temporal cost function) is expressed as the log of the temporal survival probability. And we show that the objective function of the reinforcement learning in this sense is proportional to the variational lower bound of the original problem. Finally, We empirically demonstrate that the agent learns survival behavior by using the reward function introduced in this paper.**

I. Preliminaries

Survival strategies are essential for biological agents. Many researchers have developed various types of survival agents since the early days of artificial intelligence. Ashby developed Homeostat, which dynamically stabilizes the state of the machine [1], and Walter developed simple robotic agents that can explore the environment of a room and automatically recharge their batteries at a recharging station [2]. Toda discussed the survival problem of artificial agents in the natural environment [3], [4]. He speculated about the functional requirements for an autonomous survival agent based on decision theory. Based on Toda's works, Pfeifer and Scheier pointed out the importance of 'complete' autonomous agents and research on embodied cognitive science [5]. In this sense, they developed a simple self-sufficient autonomous robot. McFarland and Bösser discussed the autonomous agent from the perspective of research on animal behavior [6]. They suggested several requirements for intelligent agents through comparison of the market economy with natural selection, and they also developed simple robots that were self-sufficient, autonomous agents [7]. Also, Lin, in a simulation study, compared several reinforcement learning (RL) architectures for a complex survival task in a non-Markovian environment [8]. Sibly and McFarland introduced the state space approach in ethology and suggested that animal behaviors are the consequence of optimal control with respect to the cost function given by the fitness function [9]. Keramati and Gutkin suggested a similar perspective, but they also suggested changes in the distance between the current homeostatic state and the optimal, desired state as the reward function of RL [10]. Konidaris and Barto developed RL architecture that automatically balances multiple required nutrients (protein, fat, water, etc.) through tuning of the reward function, which depends on the agent's homeostatic state [11]. They tested the architecture in a simulation experiment.

Ogata and Sugino developed a real robot agent intended for survival [12]. Their robot evaluates the sensor signals (motor temperature, battery level, etc.) in real time and learns to avoid undesirable stimuli. Doya and Uchibe developed robots called "Cyber rodents" intended for the study of learning agents for survival and evolutionary robotics [13]. Cyber rodents can recharge their batteries by touching special battery packs in the environment. The wireless communication modules of robots enable the software evolution of control algorithms [14].

Reward stimuli have been introduced in most of the preceding research studies, with the reward function being necessary for the RL paradigm. However, almost all of the previous studies on survival adopted hand-crafted reward functions that do not guarantee the intended behaviors, which are the survival strategies in this case. The reward and objective function given by the designer may work well in simple RL tasks. However, for more complex tasks and life-long learning settings in which agents must learn for days and months with a large amount of data, a badly hand-crafted reward function may ultimately have serious problems in terms of system performance.

Based on such considerations, we first focus on the mathematical formulation of the classical survival problem as a maximization problem of the multi-step survival probability. To solve this problem, we introduce an iterative model-based method for maximization of the objective function using an expectation-maximization (EM) algorithm with variational approximations. Surprisingly, after the M-step, the negative free-energy function (variational lower bound) in our algorithm is identical in form to the classical RL objective function with a specific reward function. Therefore, it can be maximized by classical model-free RL algorithms. Our contributions are i) a probabilistic formulation of the classical survival problem, ii) a suggested RL approach to solve this problem and iii) a demonstration that the maximization of multi-step survival probability through RL algorithms is identical to the maximization of the log of that probability from the variational lower bound.

A. Objective Functions of Reinforcement Learning Problem

Reinforcement learning (RL) is the field of research that constructs learning agents that obtain an optimal policy through interactions with the environment. We first introduce the general form of the objective function in the POMDP (partially observable Markov decision process) model. For the sake of simplicity, we restrict the discussion to a finite set of actions, states, and observations.

Many realistic environments for the agent are known to be modeled by the partially observable Markov decision process (POMDP) [15]. The POMDP model consists of the state set $\mathcal{S}$, the action set $\mathcal{A}$, the observation set $\mathcal{O}$, the transition probability to state $s' \in \mathcal{S}$ given a state $s \in \mathcal{S}$ and an action $a \in \mathcal{A}$ as $P(s'|s,a)$, the observation probability $P(o|s)$, and the reward function $r(s,a)$.

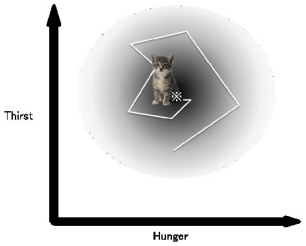

Fig. 1. The State Space Model: The figure shows the relationships between the physiological state and the viability zone. In this figure, the viability zone depends only on the two continuous variables, the "energy level" and "water level" for the sake of simplicity. The star represents the position of the optimal physiological state. The agent can remain alive in the viability zone (darker area) and should be able to recover from a position separated from this area.

At each time step t, the agent receives an observation o_t from the state s_t by the observation probability and replies with an action a_t. Then, the state of the environment changes to s_{t+1} and the agent receives a reward $r_t = r(s_t, a_t)$. In the POMDP setting, since the agent can not gain access to the true state $s \in \mathcal{S}$ of the environment, the agent needs to infer the current true state from a sequence of observations and actions. We call the sequence of observations and actions a **history** $h_t = \{o_0, a_0, o_1, a_1, \ldots, o_{t-1}, a_{t-1}, o_t\}$. Using the history, how the agent acts in the environment can be expressed by a probability $\pi(a|h)$ called the **policy**. In the POMDP model, the probability of generating the T-step **trajectory** $\tau_T = \{s_0, o_0, a_0, s_1, o_1, a_1, \ldots, s_{T-1}, o_{T-1}, a_{T-1}, s_T, o_T\}$ is

$$\begin{aligned} P(\tau_T|\pi) &= P(o_0|s_0)P(s_0) \\ &\quad \times \prod_{t=0}^{T-1} P(o_{t+1}|s_{t+1})P(s_{t+1}|s_t, a_t)\pi(a_t|h_t). \end{aligned}$$

Therefore, the expectation of the T-step average reward is given by

$$J_T(\pi) = \sum_{\tau_T} P(\tau_T|\pi)\Big[\frac{1}{T}\sum_{t=0}^{T-1} r_t\Big].$$

We denote the limit $T \to \infty$ by $J(\pi) = \lim_{T\to\infty} J_T(\pi)$. This $J_T(\pi)$ or $J(\pi)$ (or the product with some constant) is the typical objective function in the reinforcement learning literature[1]. The objective of reinforcement learning is to find the optimal policy π^* that maximally achieves the objective function, defined above, through interactions with the environment.

II. Survival Problem

In this section, we formulate the survival problem from the models of an animal proposed by Ashby [1] and a similar idea suggested by Sibly and McFarland [9] and McFarland and Böuser [6] from the view point of ethology. In their model, an animal has several variables that are observed by the animal and have some importance for sustaining life (for example, the water level and the energy level in the body, as shown in Figure 1). We call these variables the 'physiological state'. On the other hand, an agent also has a 'perceptual state', which is the perception of the environmental stimuli (vision, touch, etc.). The combined physiological state and perceptual state, which may be represented in the animal's brain, is called the 'motivational state' [6][2]. The animal has one compact manifold, which is defined in the physiological state space. This manifold is called the viability zone [16], and we define the state of the animal as 'Alive' when the current physiological state is in the manifold. The adaptive behavior is expressed as an optimal process that steers the physiological state toward the optimal state using the observed data from outside and inside the body.

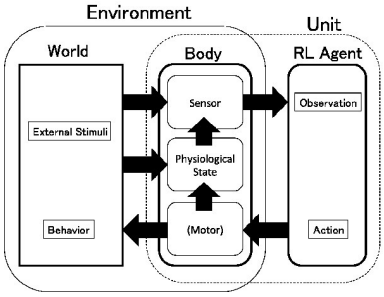

Fig. 2. The settings in the survival problem with an RL agent. A unit of the agent consists of a body and a RL agent. The RL agent interacts with the world through the body. The world receives the motor outputs from the body, returns the (external) stimulus, and changes the physiological state of the body. The body receives an external stimulus through the sensors while monitoring the physiological state through other sensors. Then, the body sends an observation to the RL agent. The RL agent receives the observation and responds with an action.

A. Formulation of the Survival Problem

The assumptions of the problem are as shown in Figure 2. Similar settings have been suggested elsewhere [1], [11], [17]–[19]. A unit of the agent consists of an RL agent and a body, and the RL agent interacts with the world through the body. Because the sensors may not perfectly determine the current situation of the world and the physiological state of the body, the sensor may exhibit only partial observability. Since sensing of the physiological state is a process inside of the body, we may be able to assume that there is no information loss with this sensing.

We now formalize the survival problem as an optimization problem. We mostly follow the usual definition of the dynamics in the POMDP model explained in the previous section. Like the POMDP model, the agent interacts with the environment. At the state $s_t \in \mathcal{S}$, the agent receives the current observation $o_t \in \mathcal{O}$ with probability $P(o|s)$. Then, the agent takes an action $a_t \in \mathcal{A}$ following some policy π; the state then changes at the next time step following the probability $P(s_{t+1}|s_t, a_t)$. Because o_t is the observation which consists of the temporal stimulus to the agent at the time step t, we can understand that the observation in POMDP is a generalization of the motivational state. In this definition, we do not explicitly separate the observation caused by external stimuli (the perceptual state; vision, touch, etc.) from the observation

[1]Even though many studies derive the algorithm from another objective function like $\sum_{t=0}^{T-1}\gamma^t r_t$, the performance of the algorithm is usually evaluated by the total or the average reward criterion.

[2]The terms 'physiological state', 'perceptual state' and 'motivational state' may be misleading in this paper, because these "states" do not necessarily follow Markovian dynamics.

caused by internal stimuli (the observed physiological state; energy level, water level, etc.).

In order to formulate the survival problem, we introduce a binary signal A_t instead of the reward, which represents the "alive flag" of the agent at times step $t = 0, 1, 2, 3 \cdots$. $A_t = 1$ represents that the agent is 'Alive' at the time step t, and $A_t = 0$ represents 'Dead'. To generalize the problem, we soften the definition of the boundary of the viability zone by introducing the temporal survival probability $P(A_{t+1} = 1|s_t) > 0$. Because the survival probability is ultimately caused by the physiological state of the agent (animal), we assume that the survival probability is only dependent on the current state. However, the following discussion is directly applicable for other definitions of the temporal survival probability, including $P(A_{t+1} = 1|o_t)$, $P(A_{t+1} = 1|s_t, a_t)$, and so on. Also, we assume $A_0 = 1$, and this probability is known for the agent unit (either in the body or RL agent).

A natural interpretation of 'survival' is to stay alive as long as possible in the environment, so the agent requires policy π that realizes the signal sequence $A_t = 1$ $(t = 0, 1, \dots)$. Using these definitions, the multi-step survival probability given a policy π is expressed by a joint probability

$$P(\bar{A}_T|\pi), \tag{1}$$

where $\bar{A}_T$ denotes the sequence $\{A_0 = 1, A_1 = 1, \dots, A_{T+1} = 1\}$. We then define the objective of the agent for the survival problem as the maximization of the probability defined by (1).

B. Maximization of Survival Probability by Variational Method

In the following discussion, we show that the maximization of the objective function (1) through maximization of the variational bound can be reduced to maximization of the conventional objective function of RL.

First, we discuss the situation of the planning problem, in which we search for the optimal policy given the true transition probability $P(s'|s, a)$, the observation probability $P(o|s)$, and the temporal survival probability $P(A|s)$. The logarithm of the objective function (1) can be transformed by introducing the arbitrary probability distribution $Q(\tau)$ on the T-step trajectory τ as

$$\begin{aligned}
\log P(\bar{A}_T|\pi) &= \sum_\tau Q(\tau) \log P(\bar{A}_T|\pi) \\
&= \sum_\tau Q(\tau) \log \frac{P(\bar{A}_T, \tau|\pi)}{P(\tau|\bar{A}_T, \pi)} \frac{Q(\tau)}{Q(\tau)} \\
&= \sum_\tau Q(\tau) \Big(\log P(\bar{A}_T|\tau) - \log \frac{Q(\tau)}{P(\tau|\pi)} \Big) \\
&\qquad + \sum_\tau Q(\tau) \log \frac{Q(\tau)}{P(\tau|\bar{A}_T, \pi)} \\
&= -F(Q(\cdot), P(\cdot|\pi)) + \mathrm{KL}(\mathrm{Q}(\cdot)||\mathrm{P}(\cdot|\bar{\mathrm{A}}_\mathrm{T}, \pi)).
\end{aligned}$$

The relationship $P(\bar{A}_T|\pi)P(\tau|\bar{A}_T, \pi) = P(\bar{A}_T, \tau|\pi)$ was used at the second equality. $P(\tau|\bar{A}_T, \pi)$ is the posterior of $P(\tau|\pi)$ given $\bar{A}$, which is $P(\tau|\bar{A}_T, \pi) = \frac{P(\bar{A}_T|\tau)P(\tau|\pi)}{\sum_{\tau'} P(\bar{A}_T|\tau')P(\tau'|\pi)}$ from Bayes's theorem, and $P(\bar{A}_T, \tau|\pi) = P(\bar{A}_T|\tau)P(\tau|\pi)$ was used at the third equality. The first term on the RHS in the last row is

$$-F(Q(\cdot), P(\cdot|\pi)) = \sum_\tau Q(\tau) \Big(\log P(\bar{A}_T|\tau) - \log \frac{Q(\tau)}{P(\tau|\pi)} \Big),$$

and F is called the free energy from the analogy with the statistical mechanics. If there is some restriction on the distribution $Q(\tau)$ (for example, $Q(\tau)$ is given by a specific class of the probability distribution), F is called the variational free energy. $\mathrm{KL}(Q(\cdot)||P(\cdot))$ is the Kullback-Leibler (KL) divergence $\mathrm{KL}(Q(\cdot)||P(\cdot)) = \sum_\tau Q(\tau) \log \frac{Q(\tau)}{P(\tau)}$. The KL divergence is known to be non-negative and zero only if the two probability distributions P and Q are equivalent. Because of the non-negativity of the KL divergence and the above equality, the log-probability $\log P(\bar{A}_T|\pi)$ is lower bounded by $-F$. Hence, this value is termed the variational (lower) bound.

Also, from the equality above

$$\log P(\bar{A}_T|\pi) + F(Q(\cdot), P(\cdot|\pi)) = \mathrm{KL}(\mathrm{Q}(\cdot)||\mathrm{P}(\cdot|\bar{\mathrm{r}}_\mathrm{T}, \pi)),$$

and there is a relationship

$$\underset{Q}{\mathrm{argmin}}[F(Q(\cdot), P(\cdot|\pi))] = \underset{Q}{\mathrm{argmin}}[\mathrm{KL}(\mathrm{Q}(\cdot)||\mathrm{P}(\cdot|\bar{\mathrm{A}}_\mathrm{T}, \pi))]$$

because $\log P(\bar{A}_T|\pi)$ is not a function of Q.

In the maximization problem of the likelihood $\log P(\bar{A}_T|\pi)$, the method that introduces the restricted class of $Q(\tau)$ and maximizes the variational bound $-F$ is known to be the variational method and it is widely used in machine learning [20]. The probability distribution of the T-step trajectory $\tau = \{s_0, o_0, a_0, s_1, o_1, a_1, \dots, s_{T-1}, o_{T-1}, a_{T-1}, s_T, o_T\}$ given a policy π is

$$\begin{aligned}
P(\tau|\pi) &= P(o_0|s_0)P(s_0) \\
&\quad \times \prod_{t=0}^{T-1} P(o_{t+1}|s_{t+1})P(s_{t+1}|s_t, a_t)\pi(a_t|h_t).
\end{aligned}$$

And we restrict the distribution $Q(\tau)$ with arbitrary policy $\pi_Q(a|h)$ to

$$\begin{aligned}
Q(\tau|\pi_Q) &= P(o_0|s_0)P(s_0) \\
&\quad \times \prod_{t=0}^{T-1} P(o_{t+1}|s_{t+1})P(s_{t+1}|s_t, a_t)\pi_Q(a_t|h_t).
\end{aligned}$$

By using these distributions, we introduce the EM algorithm as Algorithm 1. The maximization in the M-step is simply the replacement of π^k by π^k_Q from the equality

$$\begin{aligned}
&\underset{\pi}{\mathrm{argmax}}[-F(Q(\cdot|\pi_Q), P(\cdot|\pi))] \\
&= \underset{\pi}{\mathrm{argmax}} \Big[\sum_\tau Q(\tau|\pi_Q) \Big(\log P(\bar{A}_T|\tau) - \log \frac{Q(\tau|\pi_Q)}{P(\tau|\pi)} \Big) \Big] \\
&= \underset{\pi}{\mathrm{argmin}} \Big[\sum_\tau Q(\tau|\pi_Q) \log \frac{Q(\tau|\pi_Q)}{P(\tau|\pi)} \Big] \\
&= \underset{\pi}{\mathrm{argmin}}[\mathrm{KL}(Q(\cdot|\pi_Q)||P(\cdot|\pi))] \\
&= \pi_Q.
\end{aligned}$$

Because of the restriction on $Q(\tau|\pi_Q)$, the minimization of the KL divergence in the E-step is a variational sense and may

Algorithm 1

1. Set $k = 0$ and an arbitrary policy π^0.
2. (E-step) Obtain π_Q^k by optimization

$$\begin{aligned}\pi_Q^k &= \underset{\pi_Q}{\mathrm{argmin}}[\mathrm{KL}(\mathrm{Q}(\cdot|\pi_\mathrm{Q})||\mathrm{P}(\cdot|\bar{\mathrm{A}}_\mathrm{T};\pi^\mathrm{k}))] \\ &= \underset{\pi_Q}{\mathrm{argmax}}[-F(Q(\cdot|\pi_Q), P(\cdot|\pi^k))].\end{aligned}$$

3. (M-step) Update π by using π_Q^k. That is

$$\begin{aligned}\pi^{k+1} &= \underset{\pi}{\mathrm{argmax}}[-F(Q(\cdot|\pi_Q^k), P(\cdot|\pi))] \\ &= \pi_Q^k.\end{aligned}$$

4. **if** π^{k+1} is converged, **then**
5. return π^{k+1}
6. **else**
7. $k \leftarrow k + 1$ and go to E-step.
8. **end if**

not be zero. If the environment is MDP and we assume that π_Q is a Markov policy, Rawlik et al. derived the analytical solution of the E-step and it is given by the softmax policy $\pi_Q(a|s) = \exp\{\Psi(s,a)\}/Z$, where $\Psi(s,a)$ is some energy function and Z is the normalization term [21].

In POMDP, on the other hand, no analytical solution to the E-step is known. To tackle this problem, we may parametrize the policies π and π_Q by parameters θ, ϕ as π_θ, π_ϕ. And then, we assume $\theta = \phi \Rightarrow \pi_\theta = \pi_\phi$. Therefore $\theta = \phi \Rightarrow P(\tau|\pi_\theta) = Q(\tau|\pi_\phi)$ by definition. The variational method that introduces the parametrized variational distribution Q_ϕ and maximizes the variational bound $-F(\phi,\theta) := -F(Q(\cdot|\pi_\phi), P(\cdot|\pi_\theta))$ by gradient methods are well known in the neural computing community [22]–[25]. Following this idea, we replace the E-step and M-step in Algorithm 1 by

$$\phi^k = \underset{\phi}{\mathrm{argmax}}[-F(\phi, \theta^k)] \tag{2}$$

and

$$\theta^{k+1} = \underset{\theta}{\mathrm{argmax}}[-F(\phi^k, \theta)] = \phi^k. \tag{3}$$

If each stage of the algorithm is performed exactly, a monotonic increase of the variational bound

$$-F(\phi^0,\theta^1) \leq -F(\phi^1,\theta^2) \leq -F(\phi^2,\theta^3) \leq \ldots$$

after the M-step is guaranteed. Moreover, there is a relationship after the M-step

$$\begin{aligned}-F(\phi^k,\theta^{k+1}) &= -F(\theta^{k+1},\theta^{k+1}) \\ &= \sum_\tau P(\tau|\pi_{\theta^{k+1}}) \log P(\bar{A}_T|\tau) \\ &= \sum_\tau P(\tau|\pi_{\theta^{k+1}}) \Big[\sum_{t=0}^{T-1} \log P(A_{t+1}=1|s_t)\Big] \\ &= TJ_T(\pi_{\theta^{k+1}})\end{aligned}$$

in which $J_T(\pi)$ denotes the objective function of the reinforcement learning with respect to the reward function $r_t = \log P(A_{t+1} = 1|s_t)$. Here, the equality $\log P(\bar{A}_T|\tau) = \sum_{t=0}^{T-1} \log P(A_{t+1} = 1|s_t)$ is used at the third equality. Because T is a constant, an increase of the variational bound is equivalent to the increase of $J_T(\pi_{\theta^{k+1}})$. Therefore, from the discussion above, the maximization of the log-form of the objective function $\log P(\bar{A}_T|\pi_\theta)$ through the variational bound $-F(\phi,\theta)$ is reduced to the maximization of $J_T(\pi_\theta)$ with respect to the parameter of the agent θ.

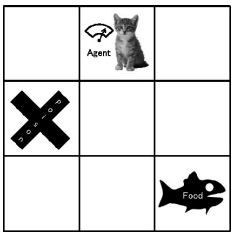

Fig. 3. The grid world environment. There are two objects, **A** and **B**, but the agent initially does not know which corresponds to the "food" object. Further details of the environment are explained in the main text.

C. Solving the Survival Problem by Reinforcement Learning Algorithms

Now we consider the survival problem in the reinforcement learning setting; that is, the maximization of the log of the objective function $\log P(\bar{A}_T|\pi)$ while the agent cannot access the true environment model $P(o|s)$, $P(s'|s,a)$. In this setting, we can not perform the iterative algorithms described above. However, from the discussion of the second algorithm (equation 2, 3), the variational bound is proportional to $J_T(\pi_\theta)$ after each M-step. Then, in order to maximize the variational bound, we can take a direct maximization of $J_T(\pi_\theta)$ with respect to θ, instead of the exact execution of the iterative algorithm. Because $J_T(\pi_\theta)$ with reward function

$$r_t = \log P(A_{t+1} = 1|s_t)$$

is the conventional objective function of the reinforcement learning paradigm, we can apply the RL algorithms to the maximization of the survival probability.

III. Experiment

In the experiment, we verify the reward setting by evaluating the finite horizon survival probability in the simple grid world domain.

a) Environment: The environment consists of a 3 x 3 grid world (Figure 3). The agent selects an action at each time step from UP, DOWN, RIGHT, LEFT and EAT. When the agent takes the action UP, DOWN, RIGHT or LEFT and if the wall is not in that direction, the agent moves one step in the selected direction. Otherwise, the agent stays at the current position. In the environment, there are two types of objects, **A** and **B**, at uniformly random positions, such that the two objects never overlap. The position of the objects changes if the agent selects the EAT action at the corresponding position. Also, the position of **B** may randomly change at every time step with a probability of 0.01.

The agent has a continuous battery level $E \in [0, 100]$ that decreases by 1% at each time step. The battery level is

recharged $+5$ if the agent selects EAT at the position of object **A**. Therefore, the object **A** corresponds to the food object.

The temporal survival probability of the agent when $A_t = 1$ is defined as $P(A_{t+1} = 1|s_t) = P(A_{t+1} = 1|E_t, C_t) = f(E_t)g(C_t)$ where E_t is the battery level at time $t = 1, 2, 3, \ldots$, $C_t \in \{0, 1\}$ is the flag bit whether the agent ate the object **B** ($C = 1$) or not ($C = 0$). f and g are defined as

$$f(E_t) = \exp\Big\{-\frac{(E_t - 60)^2}{1000}\Big\}$$

and

$$g(C_t) = \begin{cases} 0.5 & (C_t = 1) \\ 1 & (C_t = 0). \end{cases}$$

The observation of the agent is defined by the set $o = \{x, p^A, p^B, c, \hat{E}\}$, where x is the position of the agent, p^A is the position of object **A**, p^B is the position of object **B**, c_t is the type of object that the agent EATs ($c = 1$ for nothing, $c = 2$ for **A**, and $c = 3$ for **B**), and $\hat{E}$ is the discrete state of the current battery level. In this experiment, we discretized the continuous battery level $[0, 100]$ into 20 discrete regions, and $\hat{E}$ receives the class of the corresponding region of the current battery level E. Therefore, this environment is a simple POMDP setting because of the discretization of the battery level. In order to survive in this environment, the agent must take the food (**A**), avoid the poisonous object (**B**) and regulate its energy level (E). Even though this might be an over simplified model of the biological agent, this kind of situation will occur everywhere in the life of animals. Importantly, in this setting, agents initially do not know which object information (p^A, p^B) corresponds to "food" or "poison". Then the agent has to associate these objects with changes in the homeostatic values (E and c). Also, agents never receive positive rewards when they take food and the reward values for food-capture depend on the agents' battery level. Therefore, this experiment is fundamentally different from task-oriented problems like "food capturing".

b) Agent Settings: The Sarsa(λ) agent was used in this task and the action-value function in expressed by the tabular function. In this experiment, the learning rate α was 0.1, the discount rate γ, 0.95 and the decay rate of the eligibility trace λ, 0.1. The agent follows the ϵ-greedy policy in which the action is almost entirely selected by greedy action selection $a = \mathrm{argmax}_b\, Q(s, b)$ but, with a small probability of $\epsilon = 0.01$, the action is selected from the uniformly random distribution over the action set. The training procedure of the agent was as follows. An episode starts at the optimal battery level (that is, $B_t = 60$) with a random allocation of objects in the environment. At each time step, the alive flag is updated according to the temporal survival probability $P(A_{t+1}|s_t)$. If the agent receives $A_{t+1} = 0$ after the t-th update, the episode ends and the next episode starts.

c) Results: Figure 4 shows the evolution of the median survival time along the number of episodes. Evaluation was done by freezing parameters of the agent every 1000 episodes, and the agent is tested in 1000 episodes without learning. The solid line represents the median of the survival time of the Sarsa(λ) agent with the survival reward settings described above. The dashed line represents an agent that randomly

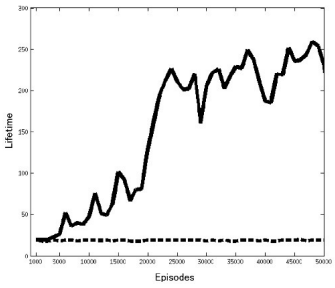

Fig. 4. The median of the survival time step along the episodes. The details are explained in the main text.

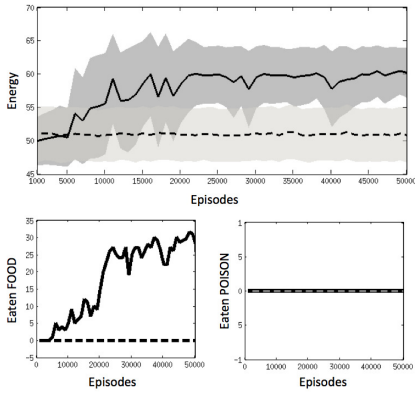

Fig. 5. Top: the evolution of the mean battery level and the standard deviation at death. Bottom left: median number of **A** (food) eaten by the agent during the evaluation step. Bottom right: the median of the number of **B** (poison) eaten by the agent during the evaluation step. Solid line: RL agent with survival reward settings. Dashed line: the agent with random action selection.

and uniformly selects one action among the 5 actions. The growth of the lifetime clearly shows that the Sarsa(λ) agent successfully learns the survival strategy during the process. The results of the random agent show that the environment has sufficient complexity that the random agent cannot stay alive longer than 25 time steps. Figure 5 shows the battery level of the agent at its death and the median amount of **A** (left panel) and **B** (right panel) consumed in the evaluation process after the corresponding episodes. From these figures, we can know that the battery level is successfully controlled around the desired level ($E = 60$) and that the amount of food eaten (**A**) increased. On the other hand, the number of poisonous objects (**B**) eaten is always zero. This result also supports the successful learning of the survival strategy by Sarsa(λ) with only the survival rewards.

IV. Discussion

In our approach, we introduced the first "fundamental" reward function of RL for the general survival problem, which so far has been only heuristically defined in previous studies. The key is to soften the definition of the viability zone with the temporal survival probability, so that the reward function is simply the log of the temporal survival probability. Using this setting, the agents can learn the survival policy with respect to the maximization of the survival probability in the future. The source of the reward function, the temporal survival

probability, has an explicit meaning and may be obtainable through the evolutionary process. However, even though our reward setting is fundamental for survival, it may not be the "optimal reward" in terms of learning efficiency for the survival policy. It is known that there are reward settings that have the same optimal policy but a different learning speed for the RL agent [26]. And, recently, the pre-training approach for the RL has been successfully applied to the real-robot domain with direct visual image inputs [27]. Therefore, to speed up learning and achieve a robotic agent that fits the unknown-dynamic environment in its (sometimes physical) lifetime, the survival problem should be examined to determine how to equip an agent with moderate prior knowledge of the environment, including the reward function and the state transition dynamics.

The relationship between the planning problem and the inference problem is a hot topic in recent machine learning communities [28]–[30]. Vlassis and Toussaint [31] introduced the perspective that model-free reinforcement learning can be treated as an application of a stochastic EM algorithm to the maximization of the mixture likelihood $p(R = 1; \pi)$. Our objective function (1) and its lower bound were briefly introduced by Toussaint [32] in the context of a stochastic control problem. In this context, the goal of our study is the maximization of the joint probability (1) from the beginning. Further, we demonstrated the relationships between its lower bound and the EM-based approach, including the POMDP case.

V. Conclusion

We have discussed the survival problem of the agent in the environment, and have shown that the survival problem can be reduced to an RL problem in (PO)MDPs with a specific reward function. Because of the popularity of the (PO)MDP assumptions in the research of autonomous agents, especially those concerning models of animals, this formulation may be seen as the basis of a truly autonomous agent for survival.

Acknowledgment

I would like to thank Makoto Otsuka and Stephany Nix for helpful comments to improve the quality of this paper.

References

[1] W. R. Ashby, *Design for a Brain*. Springer Science & Business Media, 1960.

[2] W. Walter, *The living brain.* Norton, 1953.

[3] M. Toda, "The design of a fungus-eater: A model of human behavior in an unsophisticated environment," *Behavioral Science*, vol. 7, no. 2, pp. 164–183, 1962.

[4] ——, *Man, robot, and society: Models and speculations*. M. Nijhoff Pub., 1982.

[5] R. Pfeifer and C. Scheier, *Understanding intelligence*. MIT press, 1999.

[6] D. McFarland and T. Bösser, *Intelligent behavior in animals and robots*. MIT Press, 1993.

[7] D. McFarland and E. Spier, "Basic cycles, utility and opportunism in self-sufficient robots," *Robotics and Autonomous Systems*, vol. 20, no. 2, pp. 179–190, 1997.

[8] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine learning*, vol. 8, no. 3-4, pp. 293–321, 1992.

[9] R. Sibly and D. McFarland, "On the fitness of behavior sequences," *American Naturalist*, pp. 601–617, 1976.

[10] M. Keramati and B. S. Gutkin, "A reinforcement learning theory for homeostatic regulation," in *Advances in Neural Information Processing Systems*, 2011, pp. 82–90.

[11] G. Konidaris and A. Barto, "An adaptive robot motivational system," in *From Animals to Animats 9*. Springer, 2006, pp. 346–356.

[12] T. Ogata and S. Sugano, "Emergence of robot behavior based on self-preservation. research methodology and embodiment of mechanical system," *Journal of the Robotics Society of Japan*, vol. 15, no. 5, pp. 710–721, 1997.

[13] K. Doya and E. Uchibe, "The cyber rodent project: Exploration of adaptive mechanisms for self-preservation and self-reproduction," *Adaptive Behavior*, vol. 13, no. 2, pp. 149–160, 2005.

[14] S. Elfwing, E. Uchibe, K. Doya, and H. I. Christensen, "Biologically inspired embodied evolution of survival," in *Evolutionary Computation, 2005. The 2005 IEEE Congress on*, vol. 3. IEEE, 2005, pp. 2210–2216.

[15] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *arXiv preprint cs/9605103*, 1996.

[16] J.-A. Meyer and A. Guillot, "Simulation of adaptive behavior in animats: Review and prospect," in *In J.-A. Meyer and S.W. Wilson (Eds.) From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, 1991, pp. 2–14.

[17] D. McFarland and A. Houston, *Quantitative ethology*. Pitman Advanced Pub. Program, 1981.

[18] D. Cañamero, "Modeling motivations and emotions as a basis for intelligent behavior," in *Proceedings of the first international conference on Autonomous agents*. ACM, 1997, pp. 148–155.

[19] A. G. Barto, S. Singh, and N. Chentanez, "Intrinsically motivated learning of hierarchical collections of skills," in *Proc. 3rd Int. Conf. Development Learn*, 2004, pp. 112–119.

[20] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.

[21] K. Rawlik, M. Toussaint, and S. Vijayakumar, "On stochastic optimal control and reinforcement learning by approximate inference," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 2013, pp. 3052–3056.

[22] P. Dayan and G. E. Hinton, "Varieties of helmholtz machine," *Neural Networks*, vol. 9, no. 8, pp. 1385–1403, 1996.

[23] R. Ranganath, S. Gerrish, and D. M. Blei, "Black box variational inference," *arXiv preprint arXiv:1401.0118*, 2013.

[24] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[25] A. Mnih and K. Gregor, "Neural variational inference and learning in belief networks," *arXiv preprint arXiv:1402.0030*, 2014.

[26] A. Y. Ng, D. Harada, and S. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *ICML*, vol. 99, 1999, pp. 278–287.

[27] S. Lange, M. Riedmiller, and A. Voigtlander, "Autonomous reinforcement learning on raw visual input data in a real world application," in *Neural Networks (IJCNN), The 2012 International Joint Conference on*. IEEE, 2012, pp. 1–8.

[28] M. Toussaint, S. Harmeling, and A. Storkey, "Probabilistic inference for solving (po) mdps," University of Edinburgh, Informatics Research Report 0934, 2006.

[29] E. Todorov, "General duality between optimal control and estimation," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 4286–4292.

[30] H. J. Kappen, V. Gómez, and M. Opper, "Optimal control as a graphical model inference problem," *Machine learning*, vol. 87, no. 2, pp. 159–182, 2012.

[31] N. Vlassis and M. Toussaint, "Model-free reinforcement learning as mixture learning," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 1081–1088.

[32] M. Toussaint, "Robot trajectory optimization using approximate inference," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 1049–1056.